


INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 5 : G06F 9/38	A1	(11) International Publication Number: WO 93/22722 (43) International Publication Date: 11 November 1993 (11.11.93)
(21) International Application Number: PCT/JP93/00553 (22) International Filing Date: 27 April 1993 (27.04.93) (30) Priority data: 07/877,451 1 May 1992 (01.05.92) US (71) Applicant: SEIKO EPSON CORPORATION [JP/JP]; 4-1, Nishi-Shinjuku 2-chome, Shinjuku-ku, Tokyo 163 (JP). (72) Inventors: WANG, Johannes ; 25 King Street, Redwood City, CA 94062 (US). GARG, Sanjiv ; 46820 Sentinel Drive, Fremont, CA 94539 (US). DEOSARAN, Trevor ; 555 E. Washington Ave., Sunnyvale, CA 94086 (US).		(74) Agents: SUZUKI, Kisaburo et al.; Seiko Epson Corporation, 4-1, Nishi-Shinjuku 2-chome, Shinjuku-ku, Tokyo 163 (JP). (81) Designated States: JP, KR, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i>
(54) Title: A SYSTEM AND METHOD FOR RETIRING INSTRUCTIONS IN A SUPERSCALAR MICROPROCESSOR		
(57) Abstract <p>A system and method for retiring instructions in a superscalar microprocessor which executes a program comprising a set of instructions having a predetermined program order, the retirement system for simultaneously retiring groups of instructions executed in or out of order by the microprocessor. The retirement system comprises a done block for monitoring the status of the instructions to determine which instruction or group of instructions have been executed, a retirement control block for determining whether each executed instruction is retirable, a temporary buffer for storing results of instructions executed out of program order, and a register array for storing retirable instruction results. In addition, the retirement control block further controls the retiring of a group of instructions determined to be retirable, by simultaneously transferring their results from the temporary buffer to the register array, and retires instructions executed in order by storing their results directly in the register array. The method comprises the steps of monitoring the status of the instructions to determine which group of instructions have been executed, determining whether each executed instruction is retirable, storing results of instructions executed out of program order in a temporary buffer, storing retirable instruction results in a register array and retiring a group of retirable instructions by simultaneously transferring their results from the temporary buffer to the register array, and retiring instructions executed in order by storing their results directly in the register array.</p>		



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	FR	France	MR	Mauritania
AU	Australia	GA	Gabon	MW	Malawi
BB	Barbados	GB	United Kingdom	NL	Netherlands
BE	Belgium	GN	Guinea	NO	Norway
BF	Burkina Faso	GR	Greece	NZ	New Zealand
BG	Bulgaria	HU	Hungary	PL	Poland
BJ	Benin	IE	Ireland	PT	Portugal
BR	Brazil	IT	Italy	RO	Romania
CA	Canada	JP	Japan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SK	Slovak Republic
CI	Côte d'Ivoire	LJ	Liechtenstein	SN	Senegal
CM	Cameroon	LK	Sri Lanka	SU	Soviet Union
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	MC	Monaco	TG	Togo
DE	Germany	MG	Madagascar	UA	Ukraine
DK	Denmark	ML	Mali	US	United States of America
ES	Spain	MN	Mongolia	VN	Viet Nam
FI	Finland				

DESCRIPTION

A SYSTEM AND METHOD FOR RETIRING INSTRUCTIONS IN A SUPERSCALAR MICROPROCESSOR

5

CROSS-REFERENCE TO RELATED APPLICATIONS

The following are commonly owned, co-pending applications:

10

"Superscalar RISC Instruction Scheduling", Serial No. 07/860,719, filed on March 31, 1991 (Attorney Docket No. SP035); and

"High Performance RISC Microprocessor Architecture", Serial No. 07/817,810, filed 1/8/92 (Attorney Docket No. SP015).

The disclosures of the above applications are incorporated herein by reference.

15

BACKGROUND OF THE INVENTION

1. Field of the Invention

20 The present invention relates generally to the field of superscalar processors, and more particularly, to a system and method for retiring multiple instructions executed out-of-order in a superscalar processor.

2. Discussion of Related Art

25 One method of increasing performance of microprocessor-based systems is overlapping the steps of different instructions using a technique called pipelining. In pipelining operations, various steps of instruction execution (e.g. fetch, decode and execute) are performed by independent units called pipeline stages. The steps are performed in parallel in the various pipeline stages so that the processor can handle more than one instruction at a time.

30 As a result of pipelining, processor-based systems are typically able to execute more than one instruction per clock cycle. This practice allows the rate of instruction execution to exceed the clock rate. Processors that issue, or initiate execution of, multiple independent instructions per clock cycle are known as superscalar processors. A superscalar processor reduces the average number of cycles per instruction beyond what is possible in ordinary pipelining systems.

35 In a superscalar system, the hardware can execute a small number of independent instructions in a single clock cycle. Multiple instructions can be executed in a single cycle as long as there are no data dependencies, procedural dependencies, or resource conflicts. When such dependencies or conflicts exist, only
40 the first instruction in a sequence can be executed. As a result, a plurality of functional units in a superscalar architecture can not be fully utilized.

To better utilize a superscalar architecture, processor designers have enhanced processor look-ahead capabilities; that is the ability of the processor to examine instructions beyond the current point of execution in an attempt to find independent instructions for immediate execution. For example, if an instruction
5 dependency or resource conflict inhibits instruction execution, a processor with look-ahead capabilities can look beyond the present instruction, locate an independent instruction, and execute it.

As a result, more efficient processors, when executing instructions, put less emphasis on the order in which instructions are fetched and more emphasis on the
10 order in which they are executed. As a further result, instructions are executed out of order.

For a more in-depth discussion of superscalar processors, see Johnson, *Superscalar Microprocessor Design*, Prentice Hall, Inc. (1991).

Scenarios occur whereby the execution of the instructions is interrupted or
15 altered, and the execution must be restarted in the correct order. Two such scenarios will be described.

In a first scenario, during look-ahead operations, many processor designs employ predictive techniques to predict a branch that the program is going to follow in that particular execution. In these systems, the instructions fetched and executed
20 as a result of look-ahead operations are instructions from the branch of code that was predicted. High instruction throughput is achieved by fetching and issuing instructions under the assumption that branches chosen are predicted correctly and that exceptions do not occur. This technique, known as speculative execution, allows instruction execution to proceed without waiting for the completion of previous
25 instructions. In other words, execution of the branch target instruction stream begins before it is determined whether the conditional branch will be taken.

Since the branch predication occasionally fails, the processor must provide recovery mechanisms for canceling the effects of instructions that were speculatively executed. The processor must also provide restart mechanisms to
30 reestablish the correct instruction sequence.

In a second scenario, out-of-order completion makes it difficult to deal with exceptions. Exceptions are created by instructions when the instruction cannot be properly executed by hardware alone. These exceptions are commonly handled by interrupts, permitting a software routine to correct the situation. Once the routine is
35 completed, the execution of the interrupted program must be restarted so it can continue as before the exception.

Processors contains information that must be saved for a program to be suspended and then restored for execution to continue. This information is known as the 'state' of the processor. The state information typically includes a program

counter (PC), an interrupt address register (IAR), and a program status register (PSR); the PSR contains status flags such as interrupt enable, condition codes, and so forth.

5 As program instructions are executed, the state machine is updated based on the instructions. When execution is halted and must later be restarted (*i.e.*, one of the two above scenarios occurs) the processor looks to the state machine for information on how to restart execution. In superscalar processors, recovery and restart occur frequently and must be accomplished rapidly.

10 In some conventional systems, when instructions are executed out of order, the state of the machine is updated out of order (*i.e.*, in the same order as the instructions were executed). Consequently, when the processor goes back to restart the execution, the state of the machine has to be 'undone' to put it back in a condition such that execution may begin again.

15 To understand conventional systems, it is helpful to understand some common terminology. An in-order state is made up of the most recent instruction result assignments resulting from a continuous sequence of executed instructions. Assignments made by instructions completed out-of-order where previous instruction(s) have not been completed, are not included in this state.

20 If an instruction is completed and all previous instructions have also been completed, the instruction's results can be stored in the in-order state. When instructions are stored in the in-order state, the machine never has to access results from previous instructions and the instruction is considered 'retired.'

25 A look-ahead state is made up of all future assignments, completed and uncompleted, beginning with the first uncompleted instruction. Since there are completed and uncompleted instructions, the look-ahead state contains actual as well as pending register values.

30 Finally, an architectural state is made up of the most recently completed assignment of the continuous string of completed instructions and all pending assignments to each register. Subsequent instructions executed out of order must access the architectural state to determine what state the register would be in had the instruction been executed in order.

35 One method used in conventional systems to recover from misdirected branches and exceptions is known as checkpoint repair. In checkpoint repair, the processor provides a set of logical spaces, only one of which is used for current execution. The other logical spaces contain backup copies of the in-order state, each corresponding to a previous point in execution. During execution, a checkpoint is made by copying the current architectural state to a backup space. At this time, the oldest backup state is discarded. The checkpoint is updated as instructions are executed until an in-order state is reached. If an exception occurs, all previous

instructions are allowed to execute, thus bringing the checkpoint to the in-order state.

To minimize the amount of required overhead, checkpoints are not made at every instruction. When an exception occurs, restarting is accomplished by loading the contents of the checkpointed state preceding the point of exception, and then
5 executing the instructions in order up to the point of exception. For branch misprediction recovery, checkpoints are made at every branch and contain the precise state at which to restart execution immediately.

The disadvantage of checkpoint repair is that it requires a tremendous amount of storage for the logical spaces. This storage overhead requires additional chip real
10 estate which is a valuable and limited resource in the microprocessor.

Other conventional systems use history buffers to store old states that have been superseded by new states. In this architecture, a register buffer contains the architectural state. The history buffer is a last-in first-out (LIFO) stack containing items in the in-order state superseded by look-ahead values (*i.e.*, old values that have
15 been replaced by new values), hence the term 'history.'

The current value (prior to decode) of the instruction's destination register is pushed onto the stack. The value at the bottom of the stack is discarded if its associated instruction has been completed. When an exception occurs, the processor suspends decoding and waits until all other pending instructions are completed, and
20 updates the register file accordingly. All values are then popped from the history buffer in LIFO order and written back into the register file. The register file is now at the in-order state at the point of exception.

The disadvantage associated with the history buffer technique is that several clock cycles are required to restore the in-order state.

25 Still other conventional systems use a reorder buffer managed as a first-in first-out (FIFO) queue to restart after exceptions and mispredictions. The reorder buffer contains the look-ahead state, and a register file contains the in-order state. These two can be combined to determine the architectural state. When an instruction is decoded, it is assigned an entry at the top of the reorder buffer. When
30 the instruction completes, the result value is written to the allocated entry. When the value reaches the bottom of the buffer, it is written into the register file if there are no exceptions. If the instruction is not complete when it reaches the bottom, the reorder buffer does not advance until the instruction completes. When an exception occurs, the reorder buffer is discarded and the in-order state is accessed.

35 The disadvantage of this technique is that it requires associative lookup to combine the in-order and look-ahead states. Furthermore, associative lookup is not straightforward since it must determine the most recent assignments if there is more than one assignment to a given register. This requires that the reorder buffer be implemented as a true FIFO, rather than a more simple, circularly addressed register

array.

What is needed then is a system and method for maintaining a current state of the machine and for efficiently updating system registers based on the results of instructions executed out of order. This system and method should use a minimum of chip real estate and power and should provide quick recovery of the state of the machine up to the point of an exception. Furthermore, the system should not require complex steps of associative lookup to obtain the most recent value of a register.

SUMMARY OF THE INVENTION

The present invention is a system and method for retiring instructions issued out of order in a superscalar microprocessor system. According to the technique of the present invention, results of instructions executed out of order are first stored in a temporary buffer until all previous instructions have been executed. Once all previous instructions have been executed and their results stored in order in a register array, the results of the instruction in question can be written to the register array and the instruction is considered retired.

The register array contains the current state of the machine. To maintain the integrity of register array data, only results of instructions are not written to the register array until the results of all previous instructions have been written. In this manner, the state of the machine is updated in order, and situations such as exceptions and branch mispredictions can be handled quickly and efficiently.

The present invention comprises means for assigning and writing instruction results to a temporary storage location, transferring results from temporary storage to the register array so that the register array is updated in an in-order fashion and accessing results in the register array and temporary storage for subsequent operations.

Further features and advantages of the present invention, as well as the structure and operation of various embodiments of the present invention, are described in detail below with reference to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a data path diagram of a superscalar instruction execution unit.

Fig. 2 is a block diagram illustrating the functions of the superscalar instruction execution unit.

Fig. 3 is a diagram further illustrating the instruction FIFO and the instruction window.

Fig. 4 is a diagram illustrating instruction retirement according to the present invention.

Fig. 5A shows the configuration of an instruction window.

Fig. 5B is a diagram illustrating the assignment of instruction results to

storage locations in a temporary buffer according to the present invention.

Fig. 6A is a timing diagram illustrating data writing to a register array according to the present invention.

Fig 6B is a timing diagram illustrating writing results to four register locations per clock cycle according to the present invention.

In the drawings, like reference numbers indicate identical or functionally similar elements. Additionally, the left-most digit of a reference number identifies the drawing in which the reference number first appears.

DETAILED DESCRIPTION

10

1. Overview

The present invention provides a system and a method for retiring completed instructions such that to the program it appears that the instructions are executed sequentially in the original program order. The technique of the present invention is to store all out-of-order instruction results (results of instructions not executed in the program order) in a temporary buffer until all previous instructions are complete without any exceptions. The results are then transferred from the temporary buffer to a register array which represents the official state.

When an instruction is retired, all previous instructions have been completed and the retired instruction is officially completed. When instructions are retired according to the technique of the present invention, the state of the machine is updated in order. Therefore, when an exception occurs, out-of-order execution is suspended and all uncompleted instructions prior to the exception are executed and retired. Thus, the state of the machine is up to date as of the time of the exception. When the exception is complete, out-of-order execution resumes from the point of exception. When a branch misprediction is detected, all instructions prior to the branch are executed and retired, the state of the machine is now current, and the machine can restart at that point. All results residing in the temporary buffer from instructions on the improper branch are ignored. As new instructions from the correct branch are executed, their results are written into the temporary buffer, overwriting any results obtained from the speculatively executed instruction stream.

2. Environment

Fig. 1 illustrates a block diagram of a superscalar Instruction Execution Unit (IEU) capable of out-of-order instruction issuing. Referring to Fig. 1, there are two multi-ported register files 102A, 102B which hold general purpose registers. Each register file 102 provides five read ports and two write ports. Each write port allows two writes per cycle. In general, register file 102A holds only integer data while register file 102B can hold both floating point and integer data.

Functional units 104 are provided to perform processing functions. In this example, functional units 104 are three arithmetic logic units (ALUs) 104A, a shifter 104B, a floating-point ALU 104C, and a floating-point multiplier 104D. Floating-point ALU 104C and floating-point multiplier 104D can execute both integer and floating-point operations.

Bypass multiplexers 106 allow the output of any functional unit 104 to be used as an input to any functional unit 104. This technique is used when the results of an instruction executed in one clock cycle are needed for the execution of another instruction in the next clock cycle. Using bypass multiplexers 106, the result needed can be input directly to the appropriate functional unit 104. The instruction requiring those results can be issued on that same clock cycle. Without bypass multiplexers 106, the results of the executed instruction would have to be written to register file 102 on one clock cycle and then be output to the functional unit 104 on the next clock cycle. Thus, without bypass multiplexers 106 one full clock cycle is lost. This technique, also known as forwarding, is well known in the art and is more fully described in Hennessy et al., *Computer Architecture a Quantitative Approach*, Morgan Kaufmann Publishers (1990) on pages 260-262.

Selection multiplexers 108 provide a means for selecting the results from functional units 104 to be written to register files 102.

Fig. 2 illustrates a block diagram of IEU control logic 200 for the IEU shown in Fig. 1. IEU control logic 200 includes an instruction window 202. Instruction window 202 defines the instructions which IEU control logic 200 may issue during one clock cycle. Instruction window 202 represents the bottom two locations in an instruction buffer, which is a FIFO register containing instructions to be executed. This instruction buffer is also referred to as an instruction FIFO. As instructions are completed, they are flushed out at the bottom and new instructions are dropped in at the top. The bottom location of instruction window 202 is referred to as bucket 0 and the top location of instruction window 202 is referred to as bucket 1.

When all four instructions in bucket 0 have been retired, they are flushed out of bucket 0, the instructions in bucket 1 drop into bucket 0 and a new group of four instructions drops into bucket 1. Instruction window 202 may be implemented using a variety of techniques. One such technique is fully described in the commonly owned, co-pending application titled "Superscalar Risc Instruction Scheduling" (Serial Number 07/860,719; Attorney Docket Number SP035/1397.0170000), filed March 31, 1992, the disclosure of which is incorporated herein by reference.

In the current example, instruction window 202 contains eight instructions. Therefore, IEU control logic 200 tries to issue a maximum number of instructions from among these eight during each clock cycle. Instruction decoding occurs in decoders 203. Instruction decoding is an ongoing process performed in IEU control

logic 200. Instructions must be decoded before dependency checking (discussed below), issuing and execution occur.

IEU control logic 200 also contains register renaming circuitry (RRC) 204 which performs two related functions. The first function performed is data dependency checking. Once data dependency checking is complete, RRC 204 assigns tags to each instruction which are used to track the location of instruction operands and results.

Data dependency checking logic, residing in RRC 204, is used for checking instructions for dependencies. In checking for dependencies, the data dependency checking logic looks at the various register file source and destination addresses to determine whether one or more previous instructions must be executed before a subsequent instruction may be executed. Fig. 3 further illustrates instruction window 202 and the instruction FIFO. Referring to Fig. 3, various register file source and destination addresses 302 of the instruction I0 must be checked against the source and destination addresses of all other instructions.

Referring back to Fig. 2, since instruction window 202 in this example can contain 8 instructions, the IEU can look at eight instructions for scheduling purposes. All source register addresses must be compared with all previous destination addresses. If one instruction is dependent upon completion of a previous instruction, these two instructions cannot be completed out of order. In other words, if instruction I2 requires the results of instruction I1, a dependency exists and I1 must be executed before I2. Some instructions may be long-word instructions, which require extra care when checking for dependencies. For long-word instructions, the instructions occupy two registers both of which must be checked when examining this instruction for dependencies.

An additional function performed in RRC 204 is tag assignment. Proper tag assignment is crucial to effective instruction retirement according to the present invention. Each instruction in instruction window 202 is assigned a tag based on its location in instruction window 202, and based on the results of data dependency checking discussed above. The tag assigned to each instruction indicates where in a temporary buffer that instruction's results are to be stored until that instruction is retired and whether all of the previous instructions on which that instruction is dependent have been completed. Tag assignment and the temporary buffer are discussed in more detail below.

A further function performed by IEU control logic 200 is determining which instructions are ready for issuing. An instruction issuer 208 issues instructions to the appropriate functional unit 104 for execution. Circuitry within RRC 204 determines which instructions in instruction window 202 are ready for issuing and sends a bit map to instruction issuer 208 indicating which instructions are ready for

issuing. Instruction decode logic 203 indicates the resource requirement for each instruction. Issuer 208 also receives information from functional units 104 concerning resource availability. This information is scanned by issuer 208 and an instruction is selected for issuing.

5 Instruction issuer 208 sends a control signal 209 to multiplexers 210 telling them which instruction to send to functional units 104. Instruction issuer 208 also sends a control signal 211 to multiplexer 212 configuring it to send the appropriate register address to configure the register that is to receive the results of the instruction. Depending on the availability of functional units 104, issuer 208 may
10 issue multiple instructions each clock cycle.

Referring again to Figs. 1 and 2, once an instruction is issued to functional units 104 and executed by the same, register files 102A and 102B must be updated to reflect the current state of the machine. When the machine has to 'go back' and restart an execution because of an exception or a branch misprediction, the state of
15 the machine must reflect the up-to-date state at the time the exception or branch occurred. Even when instructions are issued and executed out of order, the state of the machine must still reflect, or be recoverable to, the current state at the time of exception or branching.

The Instruction Retirement Unit (IRU) of the present invention, retires the
20 instructions as if they were executed in order. In this manner, the state of the machine is updated, in order, to the point of the most recent instruction in a sequence of completed instructions.

The present invention provides a unique system and method for retiring instructions and updating the state of the machine such that when a restart is
25 required due to an exception or a branch misprediction, the current state up to that point is recoverable without needing to wait for the register file to be rebuilt or reconstructed to negate the effects of out-of-order executions.

3. Implementations

30 Fig. 4 illustrates a high-level diagram of an Instruction Retirement Unit 400 (referred to as "IRU 400") of the present invention. IRU 400 and its functions are primarily contained within register file 102 and a retirement control block (RCB) 409. As shown in Fig. 4, the functions performed by the environment are also critical to proper instruction retirement.

35 Referring to Fig. 4, the operation of IRU 400 will now be described. As discussed in subsection 2 of this application, the instructions executed in the superscalar processor environment are executed out of order, and the out-of-order results cannot be written to the registers until all previous instructions' results are written in order. A register array 404 represents the in-order state of the machine.

The results of all instructions completed without exceptions, who also have no previous uncompleted instructions, are stored in register array 404. Once the results are stored in register array 404, the instruction responsible for those results is considered 'retired.'

5 If an instruction is completed out of order, and there are previous instructions that have not been completed, the results of that instruction are temporarily stored in a temporary buffer 403. Once all instructions previous to the instruction in question have been executed and their results transferred to register array 404, the instruction in question is retirable, and its results can be transferred from temporary
10 buffer 403 to register array 404. Once this is done, the instruction is considered retired. A retirable instruction then, is an instruction for which two conditions have been met: (1) it is completed, and (2) there are no unexecuted instructions appearing earlier in the program order.

If the results of an executed instruction are required by a subsequent
15 instruction, those results will be made available to the appropriate functional unit 104 regardless of whether they are in temporary buffer 403 or register array 404.

Referring to Figs. 1, 2, and 4, IRU 400 will be more fully described. Register file 102 includes a temporary buffer 403, a register array 404 and selection logic 408. There are two input ports 110 used to transfer results to temporary buffer 403 and
20 register array 404. Control signals (not shown) generated in IEU control logic 200 are used to select the results in selection multiplexer 108 when the results are ready to be stored in register file 102. Selection multiplexer 108 receives data from various functional units and multiplexes this data onto input ports 110.

Two input ports 110 for each register file 102 in the preferred embodiment
25 permit two simultaneous register operations to occur. Thus, input ports 110 provide two full register width data values to be written to temporary buffer 403. This also permits multiple register locations to be written in one clock cycle. The technique of writing to multiple register address locations in one clock cycle is fully described below.

30 Figs. 5A and B illustrate the allocation of temporary buffer 403. Fig. 5A shows a configuration of instruction window 202, and Fig. 5B shows an example ordering of data results in temporary buffer 403. As noted previously, there can be a maximum of eight pending instructions at any one time. Each instruction may require one or two of temporary buffer's 403 eight register locations 0 through 7, depending on
35 whether it is a regular-length or a long-word instruction.

The eight pending instructions in instruction window 202 are grouped into four pairs. The first instructions from buckets 0 and 1 (i.e. I0 and I4) are a first pair. The other pairs, I1 and I5, etc., are similarly formed. A result of I0 (I0RD) is stored in register location 0, and a result of I4 (I4RD) is stored in register location 1. If I0 is a

long-word entry, IORD, the low-word result (result of the first half of a long-word instruction) is still stored in location 0, but now the high-word result (IORD+1, from the second half of the instruction) is stored in location 1. This means that the low-word result of I4 does not have a space in temporary buffer 403, and therefore can not be issued at this time.

Tags are generated in RRC 204 and assigned to each instruction. Each tag comprises three bits, for example, to indicate addresses for writing the instruction's results in temporary buffer 403. These three bits are assigned according to the instructions' locations in instruction window 202. Table 1 illustrates a representative assignment for these three tag bits.

INSTRUCTION	TAG	LOCATION
0	000	0
1	010	2
2	100	4
3	110	6
4	001	1
5	011	3
6	101	5
7	111	7

Table 1. Tag assignment

Each location in instruction window 202 has a corresponding location in temporary buffer 403. The least significant bit indicates the bucket in instruction window 202 where the instructions originated. This bit is interpreted differently when the bucket containing the instruction changes. For example, when all four instructions of bucket 0 are retired, the instructions in bucket 1 drop into bucket 0. When this occurs the LSB (least significant bit) of the tag that previously indicated bucket 1, now indicates bucket 0. For example, in Table 1, an LSB of 1 indicates the instructions in bucket 1. When these instructions are dropped into bucket 0, the LSB will not change and an LSB of 1 will indicate bucket 0. The tag contains information on how to handle each instruction.

When the instruction is executed and its results are output from a functional unit, the tag follows. Three bits of each instruction's tag uniquely identify the register location where the results of that instruction are to be stored. This concept will now be described in more detail. A temporary write block (not shown) looks at functional units 104, the instruction results and the tags. Each functional unit 104 has 1 bit that indicates if a result is going to be output from that functional unit 104 on the

next clock cycle. The temporary write block gets the tag for each result that will be available on the next clock cycle. The temporary write block generates an address (based on the tag) where the upcoming results are to be stored in temporary buffer 403. The temporary write block addresses temporary buffer 403 via RRC 204 on the
5 next clock cycle when the results are ready at functional unit 104.

A secondary function of the tags is to check whether the results of a particular functional unit 104 can be routed directly to the operand input of a functional unit 104. This occurs when a register value is needed immediately by a functional unit 104. The results can also be stored in register array 404 or temporary buffer 403.

10 In addition, the tags indicate to the IEU when to return those results directly to bypass multiplexers 106 for immediate use by a instruction executing in the very next clock cycle. The instruction results may be sent to either the bypass multiplexers 106, register file 102, or both.

The results of all instructions executed out of order are stored first in a
15 temporary buffer 403. As discussed above, temporary buffer 403 has eight storage locations. This number corresponds to the size of instruction window 202. In the example discussed above, instruction window 202 has eight locations and thus there are up to eight pending instructions. Consequently, up to eight instruction results may need to be stored in temporary buffer 403.

20 If an instruction is completed in order, that is all previous instructions are already completed and their results written to register array 404, the results of that instruction can be written directly to register array 404. RCB 409 knows if results can go directly to register array 404. In this situation, RCB 409 sets an external write bit enabling a write operation to register array 404. Note, in the preferred
25 embodiment, the results in this situation are still written to temporary buffer 403. This is done for simplicity.

For each instruction result in temporary buffer 403, when all previous instructions are complete, without any exceptions or branch mispredictions, that result is transferred from temporary buffer 403 to a register array 404 via selection
30 logic 408. If an instruction is completed out of order and previous instructions are not all completed, the results of that instruction remain in temporary buffer 403 until all previous instructions are completed. If one or more instructions have been completed, and they are all awaiting completion of an instruction earlier in the program order, they cannot be retired. However, once this earlier instruction is
35 completed, the entire group is retireable and can be retired.

A done block 420 is an additional state machine of the processor. Done block 420 keeps track of what instructions are completed and marks these instructions 'done' using a done flag. The done block informs a retirement control block 409 which instructions are done. The retirement control block 409, containing retirement

control circuitry checks the done flags to see if all previous instructions of each pending are completed. When retirement control block 409 is informed that all instructions previous (in the program order) to the pending instruction are completed, the retirement control block 409 determines that the pending instruction is retirable.

5 Fig. 6A is a timing diagram illustrating writing to register array 404, and Fig. 6B is a timing diagram illustrating the transfer of data from temporary buffer 403 to register array 404. Referring to Figs. 4, 6A, and 6B, the technique of writing to register array 404 will be described.

Temporary buffer 403 has four output ports F, G, H, and I that are used to
10 transfer data to register array 404. Register array 404 has two input ports, A' and B', for accepting instruction results from either temporary buffer 403 or functional units 104. Write enable signals 602 and 604 enable writes to temporary buffer 403 and register array 404, respectively, as shown at 603. Although not illustrated, there are actually 2 write enable signals 604 for register array 404. One of these enable
15 signals 604 is for enabling writes to input port A', and the other is for enabling writes to input port B'. Since there are two input ports A', and B', two writes to register array 404 can occur simultaneously.

Data to be written to register array 404 can come from either temporary buffer 403 or functional units 104 (via selection multiplexer 108 and bus 411).
20 Control signal 606 is used to select the data in transfer logic 408. When control signal 606 is a logic high, for example, data is selected from temporary buffer 403. Signal 410 is the write address, dictating the location where data is to be written in either temporary buffer 403 or register array 404. Data signal 608 represents the data being transferred from temporary buffer 403 to register array 404.
25 Alternatively, data signal 608 represents data 110 from functional units 104 via selection multiplexer 108.

Register array 404 can write 4 locations in one clock cycle. Instructions per half clock cycle can be retained simultaneously. Address 410 and write enable 604 signals are asserted first, then data 608 and control signal 606 are asserted. Control
30 signal 606 is asserted as shown at 605. During the first half of the cycle, registers corresponding to instructions I0 and I1 will be updated. During the second half of the cycle, registers corresponding to I2 and I3 will be updated. If any of the results are long words, the upper half of the word will be updated during the second cycle. Thus, two results can be simultaneously transferred and two instructions can be
35 simultaneously retired in a half a clock cycle.

Referring to Fig. 6B, read addresses 612F, 612G, 612H, and 612I are available for temporary buffer 403 output ports F through I. Data 614F, 614G, 614H, and 614I is available from temporary buffer 403 at the beginning of the clock cycle, as shown at 615. Addresses 410A are generated for input port A' and 410B are

generated for input port B'. Similarly, a write enable signal 604A for input port A' and a write enable signal 604B for input port B' are generated for each half of the clock cycle. Address 410 appearing in the first half of the clock cycle, as shown at 611A and 611B, is the location to which data is written during enable signal 604 appearing in the first half, as shown as 605A and 605B. Similarly, data is written during the second half of the clock cycle to the address 410 appearing at that time, as shown at 613A and 613B. Since data is written to A' and B' simultaneously, up to four instruction results may be written to register array 404 during one clock cycle. Therefore, up to four instructions may be retired during one clock cycle.

10 Latches in control logic 408 hold the data constant until the appropriate address 410 is present and write enable signals 604 allow the data to be written.

The process of transferring a result from temporary buffer 403 to register array 404, as described above, is called retiring. When an instruction is retired, it can be considered as officially completed. All instructions previous to that instruction have been completed without branch mispredictions or exceptions and the state of the machine will never have to be redetermined prior to that point. As a result, to the program running in the processor, it appears that the instructions are updated and executed sequentially.

Since instructions are being issued and executed out of order, subsequent instructions may require access to register values in temporary buffer 403, as well as values stored in register array 404.

Read access to temporary buffer 403 and register file 404 is controlled by register renaming circuitry RRC 204. Such read access is required by instructions executing that need results of previously executed instructions. Recall from the discussion in subsection 2 above that RRC 204 performs data dependency checking. RRC 204 knows which instructions are dependent on which instructions and which instructions have been completed. RRC 204 determines if the results required by a particular instruction must be generated by a previous instruction, *i.e.* whether a dependency exists. If a dependency exists, the previous instruction must be executed first. An additional step is required, however, when a dependency exists. This step is determining where to look for the results of the instruction. Since RRC 204 knows what instructions have been completed, it also knows whether to look for the results of those instructions in temporary buffer 403 or register array 404.

RRC 204 sends a port read address 410 to register array 404 and temporary buffer 403 to read the data from the correct location onto output lines 412. One bit of read address 410 indicates whether the location is in temporary buffer 403 or register array 404.

In the preferred embodiment of the present invention, each output port A through E of temporary buffer 403 and register array 404 has its own dedicated

address line. That is, each memory location can be output to any port.

4. **Additional Features of the Invention.**

IRU 200 also informs other units when instructions are retired. IRU 200
5 informs an Instruction Fetch Unit (IFU) when it (the IRU) has changed the state of
the processor. In this manner, the IFU can maintain coherency with IEU 100. The
state information sent to the IFU is the information required to update the current
Program Counter and to request more instructions from the IFU. In the example
above, when four instructions are retired, the IFU can increment the PC by four and
10 fetch another bucket of four instructions.

An example of the IFU is disclosed in a commonly owned, copending
application Serial No. 07/817,810 titled "High Performance RISC Microprocessor
Architecture."

In addition, according to a preferred embodiment of the present invention,
15 status bits and condition codes are retired in order as well. Each of the eight
instructions in instruction window 202 has its own copy of the status bits and
condition codes. If an instruction does not affect any of the status bits, then it
propagates the status bits from the previous instruction.

When an instruction is retired, all its status bits have to be officially updated.
20 If more than one instruction is retired in one cycle, the status bits of the most recent
(in order) instruction are used for the update.

5. **Conclusion**

While various embodiments of the present invention have been described
25 above, it should be understood that they have been presented by way of example
only, and not limitation. Thus, the breadth and scope of the present invention should
not be limited by any of the above-described exemplary embodiments, but should be
defined only in accordance with the following claims and their equivalents.

CLAIMSWHAT IS CLAIMED IS:

1. An instruction retirement system of a superscalar microprocessor which executes a program comprising a set of instructions having a predetermined program order, said retirement system for simultaneously retiring groups of instructions executed in or out of order by the microprocessor, said retirement system comprising:
 - a first means for monitoring the status of the instructions to determine which instruction or group of instructions have been executed;
 - a second means, connected to said first means, for determining whether each executed instruction is retireable;
 - a temporary buffer for storing results of instructions executed out of program order;
 - a register array, coupled to said temporary buffer, for storing retireable-instruction results; and
 - a third means, coupled to said second means, said register array, and said temporary buffer,
 - for retiring a group of said instructions determined by said second means to be retireable, by simultaneously transferring their results from said temporary buffer to said register array, and
 - for retiring instructions executed in order by storing their results directly in said register array.
2. The system of claim 1, wherein said temporary buffer means further stores results of instructions completed in the program order.
3. The system of claim 1, further comprising transfer logic for selecting and latching results to be transferred from said temporary buffer to said register array.
4. The system of claim 1, wherein said temporary buffer comprises:
 - a plurality of storage locations for storing the results of a group of

instructions;

a plurality of address ports coupled to said storage locations for addressing said storage locations;

5

a plurality of input ports coupled to said storage locations for inputting results to be stored in the temporary buffer;

a first plurality of output ports coupled to said storage locations for outputting stored results to the register array; and

10

a second plurality of output ports coupled to said storage locations for outputting stored results to a plurality of functional units.

15 5. The system of claim 1, wherein the results of up to four instructions can be retired in a single clock cycle.

6. A method for retiring instructions in a superscalar microprocessor which executes a program comprising a set of instructions having a predetermined program order, said retirement system for simultaneously retiring groups of instructions executed in or out of order by the microprocessor, said method comprising the steps of:

20

monitoring the status of the instructions to determine which group of instructions have been executed;

25

determining whether each executed instruction is retireable;

storing results of instructions executed out of program order in a temporary buffer;

storing retireable-instruction results in a register array; and

retiring a group of said instructions determined in step (b) to be retireable, by simultaneously transferring their results from said temporary buffer to said register array.

30

7. The method of claim 6, wherein step (c) further comprises the step of storing the results of instructions completed in the program order in the temporary buffer.

35




FIG.—1




FIG.-2




FIG.-3

FIG. 5A





FIG.-5B




FIG.-4

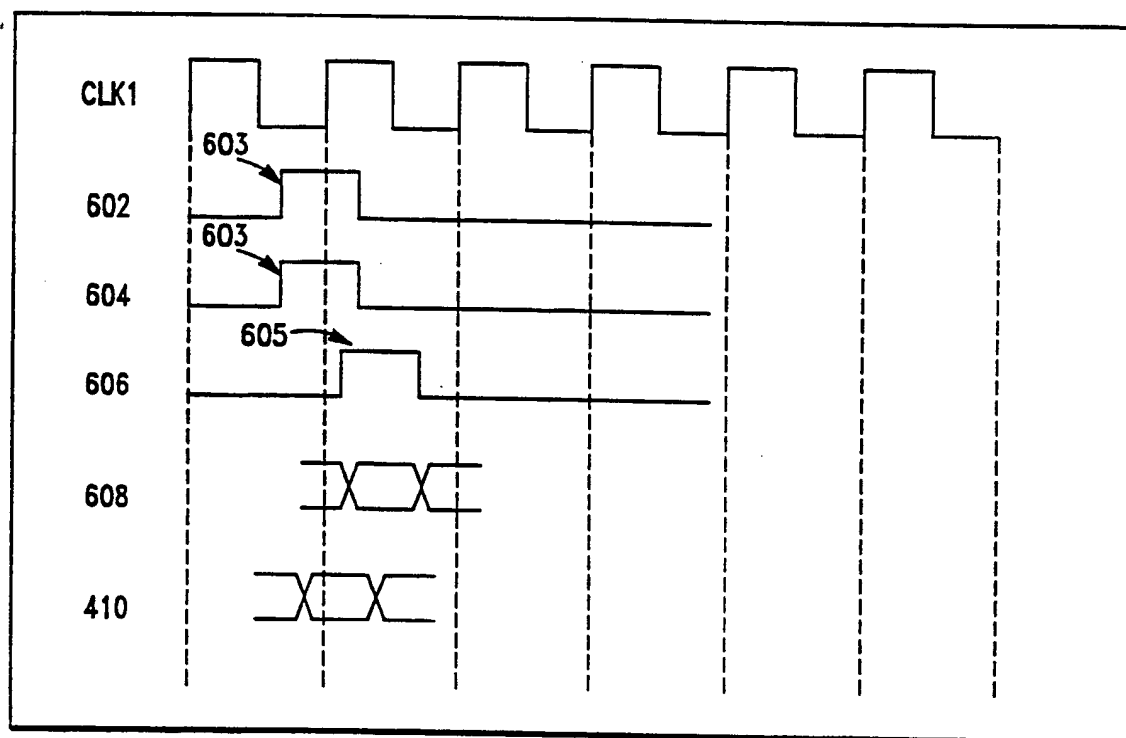


FIG.-6A




FIG.-6B

INTERNATIONAL SEARCH REPO

International Application No

PCT/JP 93/00553

I. CLASSIFICATION OF SUBJECT MATTER (If several classification symbols apply, indicate all) ⁶		
According to International Patent Classification (IPC) or to both National Classification and IPC		
Int.Cl. 5 G06F9/38		
II. FIELDS SEARCHED		
Minimum Documentation Searched ⁷		
Classification System	Classification Symbols	
Int.Cl. 5	G06F ; G11C	
Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in the Fields Searched ⁸		
III. DOCUMENTS CONSIDERED TO BE RELEVANT⁹		
Category ¹⁰	Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²	Relevant to Claim No. ¹³
X	IEEE MICRO. vol. 11, no. 3, June 1991, NEW YORK US PAGE 10-13 & 63-73 V. POPESCU ET AL. 'The Metaflow Architecture' see the whole document ---	1-3,5-7
A	WO,A,8 809 035 (ANALOG DEVICES) 17 November 1988 see abstract page 4: 'Brief Summary of the Invention' --- -/--	4
<p>¹⁰ Special categories of cited documents:</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier document but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> <p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.</p> <p>"A" document member of the same patent family</p>		
IV. CERTIFICATION		
Date of the Actual Completion of the International Search		Date of Mailing of this International Search Report
16 JULY 1993		26. 07. 93
International Searching Authority EUROPEAN PATENT OFFICE		Signature of Authorized Officer DASKALAKIS T.

Form PCT/ISA/210 (second sheet) (January 1985)

III. DOCUMENTS CONSIDERED TO BE RELEVANT (CONTINUED FROM THE SECOND SHEET)		
Category *	Citation of Document, with indication, where appropriate, of the relevant passages	Relevant to Claim No.
A	PROCEEDINGS OF THE 18TH ANNUAL INTERNATIONAL SYMPOSIUM ON COMPUTER ARCHITECTURE May 1991, TORONTO CANADA pages 276 - 286 M. BUTLER ET AL. 'Single Instruction Stream Parallelism Is Greater Than Two' see page 279, right column, paragraph 5 see page 283, left column, paragraph 5 - right column, paragraph 1 ---	1
A	M. JOHNSON 'Superscalar Microprocessor Design' 1991, PRENTICE HALL, ENGLEWOOD CLIFFS, NEW JERSEY cited in the application page 92, section 5.1.4: 'Reorder Buffer' -----	1,4

JP 9300553
SA 73159

16/07/93

EPO FORM P0479

BNSDOCID: <WO__9322722A1_I_>